

Heterogeneity-Aware Orchestration for Satellite Federated Learning

Bohang Jiang, Chao Chen, *Member, IEEE*, In-Sop Cho, Min-Su Shin,
Inkyu Lee, *Fellow, IEEE*, Seung Jun Baek, *Member, IEEE*

Abstract—The advances in non-terrestrial networks (NTNs) have transformed satellite constellations into intelligent platforms with sensing and computing capabilities, driving the emergence of satellite federated learning (FL). However, satellite FL performance is hindered by heterogeneity in communication and computation latencies of satellites, incurred by transitions into and out of low-power modes during sunlight-eclipse phases, gravitational perturbations causing antenna misalignment, or heterogeneous onboard hardware. To address these issues, this article proposes a heterogeneity-aware satellite FL system jointly optimizing the intra-orbit and inter-orbit coordinations. For intra-orbit coordination, we present Minimum Per-Satellite Latency Load (MPSLL) algorithm, which dynamically selects participants and optimizes paths for model distribution and aggregation. For inter-orbit coordination, we introduce Replacement-based Federated Learning (RFL), which retains the most recent model from each orbit and performs timely staleness-aware global aggregation. Simulations demonstrate that the proposed satellite FL system substantially outperforms baseline methods in learning performance and communication efficiency under various network heterogeneities.

Index Terms—Non-terrestrial Networks, Satellite Federated Learning, Network Heterogeneities.

I. INTRODUCTION

Recent advancements in fifth generation (5G) and ongoing sixth generation (6G) developments have accelerated the integration of satellite constellations into global communication infrastructures. Satellite networks are evolving into intelligent systems with distributed sensing, computing, and coordination capabilities.

Federated learning (FL) has emerged as a decentralized paradigm enabling collaborative model training without transmitting raw data. Nevertheless, FL performance is significantly constrained by two fundamental challenges: heterogeneity and asynchrony. Heterogeneity spans two dimensions: *statistical heterogeneity*, arising from non-identically distributed data, and *system heterogeneity*, stemming from variability in computational and communication capabilities. FedProx [1] has been proposed to tackle heterogeneity in both dimensions, and

FedAsync [2] and FedBuff [3] have been proposed to handle asynchronous client updates. However, these methods do not account for the unique constraints of satellite constellations. Satellite FL adapts this decentralized paradigm to satellite constellations and employs low Earth orbit (LEO) satellites as distributed clients and ground stations (GSs) as coordinating servers [4]. Nevertheless, satellite constellations introduce unique aspects of system heterogeneity that are not fully addressed by these general-purpose methods. Unlike terrestrial networks with static links and stable power, satellite constellations face highly dynamic constraints. For example, cyclical transitions between the sunlight and eclipse phases force energy-harvesting satellites into power-saving modes during eclipses, imposing computational constraints [5]. Gravitational perturbations from Earth's oblateness cause antenna pointing deviations, manifesting as stochastic inter-satellite link (ISL) capacity fluctuations [6]. In addition, the hardware heterogeneity is inherent in multi-operator mega-constellations. Existing approaches like FedISL [4], [7] and FedLEO [8] assume constellation homogeneity, neglecting dynamic system heterogeneity. Energy-Aware FL [9] addresses eclipse-induced energy constraints through computation scheduling. FedSN [10] employs a sub-structure scheme to accommodate heterogeneous computing and memory capabilities. However, these methods overlook the integrated latency impact of concurrent heterogeneity sources. Such oversights incur synchronization overhead and model staleness, which deteriorate the convergence rate and learning efficiency of satellite FL.

To address these challenges, this article presents a heterogeneity-aware satellite FL system that jointly optimizes intra-orbit and inter-orbit coordination, e.g., see Fig. 1 for a system overview. Specifically, intra-orbit coordination manages model training within each orbit, while inter-orbit coordination handles the global aggregation across orbits via GSs. For intra-orbit coordination, we propose the Minimum Per-Satellite Latency Load (MPSLL) orchestration algorithm that adaptively selects satellites to optimize model transmission paths, minimizing the average latency burden per selected satellite. Complementing this at the inter-orbit coordination, Replacement-based Federated Learning (RFL) is introduced as an asynchronous mechanism utilizing orbit-dedicated buffers to preserve the most recent model from each orbit, enabling fair and staleness-aware global updates. The main contributions of this article are summarized as follows:

- **Comprehensive Analysis of System Heterogeneity:** This work provides an in-depth analysis of the physical

The corresponding authors are Seung Jun Baek, Inkyu Lee, and Chao Chen. Bohang Jiang and Seung Jun Baek are with the Department of Computer Science and Engineering, College of Informatics, Korea University, Seoul, Korea (email: bohajang@korea.ac.kr; sjbaek@korea.ac.kr).

Chao Chen is with the School of Information and Electronic Engineering (Sussex Artificial Intelligence Institute), Zhejiang Gongshang University, Hangzhou, China (email: eckio_491@zjgsu.edu.cn).

In-Sop Cho and Min-Su Shin are with the Satellite Communication Infra Research Section, Electronics and Telecommunications Research Institute, Daejeon, Korea (email: lookatstar@etri.re.kr; msshin@etri.re.kr).

Inkyu Lee is with the School of Electrical Engineering, Korea University, Seoul, Korea (email: inkyu@korea.ac.kr).




Fig. 1: Overview of satellite federated learning and network system heterogeneities.

sources of heterogeneity in satellite FL, explicitly modeling how sunlight-eclipse transitions, gravitational perturbations, and hardware diversity impact both computation and communication latencies.

- **Heterogeneity-Aware Intra-Orbit Coordination:** The MPSLL orchestration algorithm is proposed to adaptively select satellites and optimize distribution and aggregation trees. By minimizing the latency burden per participant, our strategy effectively resolves synchronization bottlenecks induced by the aforementioned system heterogeneities.
- **Staleness-Aware Inter-Orbit Coordination:** RFL is developed as an asynchronous FL that utilizes dedicated buffers to address the asynchronous connectivity of multi-orbit satellite constellations. This approach enables fair and staleness-aware global updates, ensuring robust convergence even under highly irregular GS visibility.

II. FOUNDATIONS AND CHALLENGES OF SATELLITE FEDERATED LEARNING

A. System Overview of Satellite Federated Learning

Non-terrestrial networks (NTNs) are expected to play a central role in the global connectivity in 6G. In NTNs, satellite FL offers a decentralized framework for onboard training of models across satellite constellations. In a satellite FL system, satellites act as clients, while GSs serve as parameter servers (PSs). We assume that a single GS manages multiple orbits and handles the distribution, update, and aggregation of the

global model across orbits. All satellites within each orbit and GSs are assumed to be trusted entities for secure weight exchange. We assume direct communication is supported only between satellites within the same orbit via intra-orbit ISLs, not across different orbits. This is a practically reasonable assumption; the relative motion across orbits induces frequent handovers and significant Doppler shifts, making it difficult to establish stable inter-orbit satellite links [4]. As a result, the interaction among orbits must be mediated by GSs. This constraint naturally necessitates a hierarchical coordination framework across two levels: intra-orbit coordination, which includes model distribution, local training, and model aggregation within each orbit; and inter-orbit coordination, where GSs collect aggregated models from different orbits and update the global model accordingly.

Similar to typical satellite FL systems, the proposed intra-orbit coordination consists of the following operational steps:

1. **Global Model Downloading:** Each orbit begins with a stand-by mode in which satellites await an available communication window to a GS. Once the ground-satellite link becomes available, the connected satellite downloads the latest global model, which will be used as the initial model for all participating satellites in orbit.
2. **Intra-Orbit Model Distribution and Local Training:** The satellite that receives the global model initiates the distribution of the model within orbit via intra-orbit ISLs by forwarding the model to neighboring satellites until all participants have received it. Upon receiving the model, each satellite performs local training.

3. **Intra-Orbit Model Aggregation:** After completing local training, the participating satellites direct their updates toward a designated sink satellite. This forms two aggregation paths that originate from the outermost satellites and terminate at the sink. The aggregation initiates from these outermost satellites. Along these paths, intermediate satellites combine their local models with received upstream models and relay to the next node. Once the aggregated model is constructed at the sink satellite, it enters stand-by mode, awaiting the communication window to the GS to upload the model.
4. **Aggregated Model Uploading:** When the communication window opens and a ground-satellite link to the GS is established, the sink satellite uploads the aggregated model to the GS.

The effectiveness of intra-orbit coordination depends on timely model distribution, local training, and aggregation. In satellite FL, latency in these stages affects both the training efficiency within each orbit and the consistency of global model updates. Specifically, excessive latencies can cause model staleness, where outdated local updates are aggregated into the global model, degrading the convergence rate and overall model performance. We identify three key latency components in satellite FL. Firstly, the **ISL communication latency** is incurred when transmitting models via intra-orbit ISLs for model distribution and aggregation. It is determined by the achievable transmission rate, which depends on the signal-to-noise ratio (SNR), transmit power, and antenna gain. Secondly, the **onboard computation latency** arises from local model training performed by each satellite. It depends on the satellite's onboard computing power and the number of training epochs. Thirdly, the **stand-by latency** occurs when a satellite waits for a GS communication window to download the global model or upload the aggregated model. It depends on orbital geometry, link availability, and GS visibility schedules.

Next, the GS performs inter-orbit coordination to aggregate the models trained from multiple orbits and update the global model accordingly. The arrival times of trained models at the GS vary across orbits due to variations in the latencies of intra-orbit coordination and the communication window among orbits. To handle irregular arrivals of models from orbits, the PS may adopt one of the following methods for inter-orbit coordination: synchronous FL, asynchronous FL [2], or semi-asynchronous FL [3]. Later, we will propose an asynchronous FL for efficient inter-orbit coordination.

B. Challenges: System Heterogeneity in Satellite Federated Learning

As illustrated in Fig. 1, system heterogeneity in satellite FL arises mainly from three sources: sunlight-eclipse phase transitions, gravity-induced antenna pointing errors, and hardware heterogeneity. These factors cause variations in ISL quality and computing power, leading to latency fluctuations during intra-orbit coordination, as detailed below.

- **Sunlight-Eclipse Phase Transition:** Satellites periodically enter regions of sunlight and Earth's shadow during their orbital revolution around Earth. When exposed to

sunlight, satellites are in the sunlight phase, during which they operate at full power and recharge their batteries. In contrast, when passing through the shadow of Earth, satellites enter the eclipse phase, suspending energy harvesting and reducing activity to conserve power [5]. This imbalance restricts energy-intensive tasks on satellites in the eclipse phase, resulting in disparities in computation and communication capabilities. Consequently, both onboard computation latency and ISL communication latency fluctuate due to limited computing power and reduced transmission rates.

- **Gravity-Induced Pointing Error:** Gravitational perturbations from Earth's non-uniform mass distribution cause continuous shifts in satellite orientation, leading to angular misalignments between transmitting and receiving antennas. This beam pointing error, modeled as a Gaussian-distributed deviation [6], reduces antenna gain and degrades ISL quality, especially on high-frequency radio frequency (RF) links. As a result, reduced signal quality lowers the achievable transmission rate, increasing ISL communication latency during intra-orbit coordination.
- **Hardware Heterogeneity:** Satellites from different operators and technological generations exhibit significant variation in computing power, antenna gain, and transmit power [11]. Satellites with limited hardware capacities, referred to as "low-resource" nodes, typically suffer from lower computing powers and reduced transmission rates, thereby increasing onboard computation latency and ISL communication latency.

These factors incur synchronization overhead and model staleness. Furthermore, asynchronous visibility between satellites and GSs across different orbits causes irregular updates of the global model. Therefore, satellite FL systems should jointly account for dynamic intra-orbit conditions and asynchronous inter-orbit updates.

III. EFFICIENT ORCHESTRATION FOR SATELLITE FEDERATED LEARNING

In the following, we present our satellite FL system with an intra-orbit coordination for dynamic participant selection, and a replacement-based inter-orbit coordination for asynchronous and staleness-aware aggregation.

A. Intra-Orbit Coordination with Dynamic Participant Selection


The intra-orbit coordination begins with the global model downloading via connected ground-satellite links. A designated source node first downloads the global model from the GS and distributes it along a structured *distribution tree* in the orbit. As illustrated in Fig. 2(a), the distribution tree is rooted at the source node and comprises two paths: a left branch and a right branch, each terminating at a leaf node, forming a linear topology in the orbit. The nodes in the distribution tree perform local training, and the same set of satellites is reorganized into an *aggregation tree* rooted at a sink node. The sink node collects locally updated models and uploads the aggregated model to the GS.

Input: Number of satellites in each orbit K
Output: Configuration i^* with minimum ρ_i
Initialize $\rho_{\min} \leftarrow \infty$, $i^* \leftarrow \emptyset$
for each candidate source, distribution tree, and sink **do**
 if $N_i = K$ **then**
 Determine all possible aggregation trees
 else
 Select identical leaf nodes for aggregation
 for each candidate aggregation tree **do**
 Calculate total latency T_i and PSLL $\rho_i = T_i/N_i$
 if $\rho_i < \rho_{\min}$ **then**
 $\rho_{\min} \leftarrow \rho_i$ and record $i^* \leftarrow i$
return i^*


Algorithm 1: Minimum PSLL Algorithm.

We consider the problem of constructing the distribution and aggregation trees for efficient satellite FL. Previous approaches typically employ symmetric trees in which all nodes participate in the learning process as in Fig. 2(a), which may not account for variations in the node capabilities. To address this limitation, we introduce a dynamic and asymmetric tree structure in which distribution and aggregation trees are adaptively constructed according to the heterogeneity of participating satellites. As shown in Fig. 2(b), satellites with limited communication or computation capabilities, referred to as “low-resource” nodes, may be temporarily excluded from intra-orbit coordination if their participation significantly increases the latency. The “low-resource” nodes remain eligible for future iterations when their involvement becomes more favorable. This dynamic and adaptive mechanism mitigates the negative effects of satellite heterogeneity and aligns with the iterative nature of FL, where full participation in every iteration is not required.

With the introduction of the dynamic and asymmetric tree structure, the distribution tree is constructed by selecting a source node and two leaf nodes. Leaf nodes are no longer required to be adjacent. Similarly, the aggregation tree is formed by selecting a sink node and two leaf nodes. The participating nodes in the distribution and aggregation trees are identical. If all satellites in the orbit participate, we additionally select the leaf nodes of the aggregation tree independently of those of the distribution tree. Thus, the tree-construction problem reduces to selecting a configuration including: (i) source, (ii) sink, and (iii) two leaf nodes (or four when all satellites participate). We consider selecting a good configuration for dynamic tree construction, where candidate configurations must satisfy two eligibility requirements for the source and sink nodes. First, a satellite qualifies as a source node if it can download the global model immediately or within an upcoming communication window. Second, the sink node must be among the selected participants and must be able to upload the aggregated model either within its remaining communication window or during an upcoming communication window, based on its predicted visibility to the GS. The goal is to select the best configuration (source, sink, and leaf nodes) for the intra-orbit coordination. The selection criterion is based on the proposed metric called



(a) Example of symmetric distribution tree.



(b) Example of asymmetric distribution tree with dynamic participant selection.

Fig. 2: Examples of symmetric and asymmetric distribution trees.

per-satellite latency load (PSLL) as follows. Let us index the configurations with the integer i . The PSLL of configuration i , denoted by ρ_i , is defined as

$$\rho_i = \frac{T_i}{N_i}, \quad (1)$$

where T_i denotes the total learning latency (explained below), and N_i denotes the number of participating satellites. Thus, ρ_i represents the average latency burden per participating satellite for configuration i . Algorithm 1 summarizes the proposed Minimum PSLL (MPSLL) algorithm. At the beginning of each training iteration, MPSLL selects the configuration that minimizes ρ_i . MPSLL first identifies M candidate source nodes based on GS visibility. For each candidate source node, three components are selected to complete a configuration: (i) distribution leaf nodes, (ii) sink node, and (iii) aggregation leaf nodes. Each component is independently selected from all K possible choices in the worst case (all nodes participate). MPSLL exhaustively evaluates all configurations for every candidate source node. Thus, the computational complexity is given by $O(MK^3)$. Since K typically ranges from 8 to 40




Fig. 3: Overview of intra-orbit coordination with MPSLL. ① Construction of distribution tree and model download to source. ② Model distribution and local training. ③ Sink selection and model aggregation to sink. ④ Model upload to GS.

in the existing satellite FL approaches [4], [8], the complexity of MPSLL is reasonable for real-time orchestration.

The total learning latency T_i is the sum of three latency components, each defined in Section II-A. The first is the **total ISL communication latency**. It consists of the latency of the global model distribution from the source node to the satellites in the distribution tree, plus the latency of the aggregation of trained models along the aggregation tree to the sink node. The second is the **maximum onboard computation latency**. It equals the maximum of local training latencies among all participants. The aggregation is initiated after all participants have completed training, thus accounting for the synchronization overhead prior to aggregation. The third is the **stand-by latency**. It is calculated based on the visibility window of the sink node to the GS, which is predicted based on its orbital trajectory. This latency reflects the waiting time before the aggregated model can be uploaded. The ISL communication latency and onboard computation latency are estimated based on the energy availability during sunlight or eclipse phases, antenna pointing deviations, and hardware capability. Specifically, ISL communication latency is computed based on the achievable transmission rate, which is influenced by antenna misalignment that reduces antenna gain, and by reduced transmit power in eclipse phase due to limited energy. Computation latency is estimated according to

the computing power of each satellite and its state of energy availability. For example, satellites in sunlight phase operate at full computation capacity, whereas those in eclipse phase reduce computing power to conserve energy. The effective computation latency is set to the maximum among selected participants to ensure synchronization before aggregation.

Why Minimum PSLL? We explain the key idea behind MPSLL and the proposed metric ρ_i in (1) as follows. By definition, ρ_i measures the combined latencies of communication and computation per node in a distribution/aggregation tree. Consider two main components of communication and computation latencies in T_i : (i) total ISL communication latency that is proportional to the number of nodes N_i ; (ii) maximum onboard computation latency, however, takes the maximum latency among participants, which increases slower than the linear rate in N_i . Thus, T_i increases sub-linearly in the number of nodes N_i . If the system is homogeneous, i.e., if all nodes and links have identical communication/computation latencies, ρ_i will decrease in N_i . Thus, MPSLL will select the configuration with the maximum number of nodes N_i , i.e., all nodes will become participants, which reduces to the traditional node selection. Otherwise, if there exist nodes that incur excessive latency, i.e., “low-resource” nodes, configurations including those nodes will have high T_i and are likely to be excluded from the current selection. While minimizing ρ_i may




Fig. 4: Overview of Replacement-based Federated Learning (RFL).

temporarily reduce participant count, orbital periodicity ensures that excluded nodes rejoin in subsequent iterations upon transitioning into the sunlight phase or attaining favorable channel conditions. By balancing model quality with update frequency, MPSLL guarantees robust global convergence. The overall intra-orbit coordination workflow is illustrated in Fig. 3.

B. Staleness-Aware Inter-Orbit Coordination

Next, we propose an algorithm for inter-orbit coordination to aggregate the models trained through intra-orbit coordination. Since models are trained separately per orbit, asynchronous FL [2] appears efficient as it enables continuous updates without synchronous coordination. However, applying asynchronous FL directly to satellite FL poses challenges due to the system heterogeneity of satellite communication. One of the challenges is the interaction imbalance across orbits as follows. Owing to orbital geometry and GS visibility constraints, certain orbits establish connections with GS more frequently than others. Thus, if new updates are available from frequently connected orbits, they are immediately integrated into the global model. These frequently connected orbits dominate the global aggregation, introducing biased model updates toward their local data. This is particularly problematic in satellite FL; the contributions from rarely connected or delayed orbits can be diluted, leading to an *amnesia* effect, where the insights from infrequent participants are prematurely diminished from the global model, compromising both model diversity and long-term learning stability.

To address the issue, this article proposes RFL (Replacement-based Federated Learning). The key idea is that the GS maintains an *orbit-dedicated aggregation buffer*

and allocates a cell in the aggregation buffer for each orbit. Each cell stores the most recent aggregated model uploaded from its corresponding orbit. This ensures that all orbits are consistently represented in the global update process despite the variation in interaction frequency among orbits. As shown in Fig. 4, the RFL operates as follows:

1. **Orbit-Level Replacement:** When the GS receives a newly aggregated model from an orbit trained through intra-orbit coordination, it immediately replaces the previously stored model in the corresponding cell with the new one. This ensures that each orbit maintains a consistent and up-to-date representation in the global model.
2. **Staleness-Aware Aggregation:** If any cell in the buffer is updated, the GS performs global aggregation using staleness-aware aggregation. Let τ_p denote the staleness of the model from orbit p , defined as the version difference between the model used for local training and the current global model version. Let α be a positive hyper-parameter controlling the relationship between the aggregation weight and staleness. The global model w^g is updated by:

$$w^g \leftarrow \frac{1}{\beta} \sum_{p \in \mathcal{P}} \frac{w_p}{(1 + \tau_p)^\alpha},$$

where $\beta = \sum_{p' \in \mathcal{P}} (1 + \tau_{p'})^{-\alpha}$ is the normalization factor, and $\mathcal{P}$ denotes the set of all orbits.

RFL ensures that the most up-to-date model from each orbit is reflected in the global model, but balances the heterogeneous contact frequencies between the GS and orbits with staleness-aware aggregation.

C. Comparison with Existing Methodology

We qualitatively compare our satellite FL (MPSLL & RFL algorithms) with existing methodology.

1) *Intra-Orbit Efficiency*: Existing methodology lacks dynamic optimization in participant selection and distribution and aggregation tree construction. For instance, FedSat [12] ignores intra-orbit ISLs, relying instead on incremental updates. FedISL (asynchronous [4] & synchronous [7]) exploits ISLs and utilizes estimated latency for source-sink selection; however, it relies on fixed symmetric distribution and aggregation trees. FedLEO [8] optimizes source-sink selection to minimize latency. However, these methods require full participation, incurring synchronization overhead from the slowest straggler node.

2) *Inter-Orbit Robustness*: Traditional methodology fails to address system heterogeneity and interaction imbalance. Synchronous approaches, such as FedLEO [8] and FedISL (synchronous) [7], incur inefficiency from waiting for the slowest orbit. Standard asynchronous methods eliminate synchronization overhead but allow frequently connected orbits to dominate aggregation. For instance, FedSat [12] utilizes incremental updates, FedAsync (linear) [2] applies staleness-aware weighting, and FedISL (asynchronous) [4] performs data-size weighted aggregation. This introduces biased updates and causes the amnesia effect, where insights from rarely connected or delayed orbits are prematurely diluted. Furthermore, buffered approaches like FedBuff [3] (count-based trigger) and FedFa [13] (sliding-window) lack spatial fairness, risking the overwriting of models from infrequent orbits with updates from those with higher interaction frequencies. In contrast, MPSLL and RFL jointly address these limitations and guarantee learning performance in heterogeneous satellite networks.

IV. SYSTEM PERFORMANCE EVALUATION WITH NUMERICAL SIMULATION

A. Simulation Settings

LEO Satellite Constellation: A Walker-Star constellation at 500 km altitude and 50° inclination with 60 LEO satellites across 5 orbits is considered. Two GSs with PSs are located in Hangzhou, China, and Berlin, Germany. Each GS sensor has a maximum detection range of 1,000 km and a 60° half-cone angle. The constellation is analyzed using Systems Tool Kit (STK), which simulates satellite motions, GS-satellite connectivity, and sunlight-eclipse transitions [11]. The simulation spans 96 hours.

Models and Datasets: We adopt a convolutional neural network (CNN) with 3 convolutional layers (32, 64, 128 channels, kernel size 3×3 , padding 1) and 2 fully connected layers with 512 and 10 neurons. The CNN model has 1.15M parameters, and its size is 9.18MB. Local training is performed on EuroSAT [14] (10 classes, originally 64×64 , resized to 32×32 , with a random 80/20 train/test split applied over the entire dataset). For the settings of iid data, the dataset is uniformly distributed across all satellites. For the non-iid settings, a Dirichlet distribution with a concentration parameter set to 0.5 is employed to partition data across orbits, while

maintaining a uniform distribution within each orbit. The number of training epochs is set to 3, the learning rate is set to 0.01, and the batch size is set to 8.

Hardware Parameters: To account for heterogeneity in the satellite hardware, the transmit power is selected randomly from [37, 43] dBm and the antenna gain is chosen randomly from [4.98, 8.98] dBi. The latency of the global model download and upload is set to 50 seconds.

Sunlight-Eclipse Phase-Specific Parameters: We assign phase-specific parameters to satellites based on their illumination condition in each training iteration. Satellites in the sunlight phase are assumed to operate with full energy availability and use nominal transmit power for communication. Satellites in the eclipse phase lower their transmit power to 37 dBm to conserve energy. For local training, the onboard computation latency is set to 360 seconds for satellites in sunlight and 600 seconds for those in eclipse, based on the configuration in [4]. To ensure synchronization among participating satellites, the local training latency for each iteration is determined by the slowest participant; specifically, if any satellite is in eclipse, all participants adopt a latency of 600 seconds for that iteration.

Channel Condition Parameters: To model intra-orbit ISLs, we adopt a complex Gaussian channel with free-space path loss (FSPL) and gravity-induced pointing errors from orbital perturbations [6], [15]. We use a fixed carrier frequency of 2.4 GHz, and the speed of light is set to 3×10^8 m/s. Beam pointing errors are sampled for each satellite at the beginning of each iteration from a normal distribution with mean 0 and standard deviation 0.05. The resulting gain attenuation is modeled using a 3 dB beamwidth (adjusted to 0.174 radians). Noise power is computed using the Boltzmann constant of 1.380649×10^{-23} J/K, a receiver noise temperature of 354.81 K, and a channel bandwidth of 2.1 GHz. The effective SNR is calculated based on an FSPL model augmented with a gain degradation factor [6], and depends on the transmit power, antenna gains, and Euclidean distance between satellites. The resulting SNR determines the achievable transmission rate according to the Shannon capacity.

Comprehensive simulations are performed to compare the proposed satellite FL system (MPSLL & RFL algorithms) with existing satellite FL baselines, detailed in Section III-C. Gradient sparsification is used in [4], but is omitted here for a fair comparison. Additionally, FedAsync (without ISL) is implemented as a baseline that performs asynchronous aggregation without intra-orbit ISLs. The buffer size is set to 5 for both FedFa and FedBuff. We evaluate the test accuracy of the aforementioned satellite FL systems under an identical simulated wall-clock time. Python is used to partition datasets, simulate the intra-orbit ISLs model, and initialize satellite hardware capacity, while PyTorch is utilized to implement CNN model training. Each system is evaluated with random seeds ranging from 1 to 10. Linear interpolation is applied to integrate the simulation results across different schemes and seeds. In our simulation, the α in RFL is set to 3.

B. Results

As illustrated in Fig. 5(a) and Fig. 5(b), the proposed satellite FL system consistently outperforms all baselines under




Fig. 5: Performance comparison of satellite FL methods under various GS configurations, and ablation study of intra-orbit and inter-orbit coordination. Lines represent the mean accuracy across 10 independent seeds, and the shaded regions indicate the standard deviation.

both iid and non-iid settings. Specifically, MPSLL enhances participant selection by accounting for heterogeneity, while RFL enables fully asynchronous inter-orbit communication and alleviates model staleness. Their integration significantly improves adaptability to dynamic orbital conditions and ensures stable learning performance even in the presence of “low-resource” satellites. This improvement is further confirmed under a more challenging one-GS (Berlin) configuration, as shown in Fig. 5(c) and Fig. 5(d). Reducing the number of GSs from two to one introduces more severe interaction imbalance across orbits. However, the proposed satellite FL system maintains a clear performance advantage over all baselines.

An ablation study is conducted to assess the individual contributions of intra-orbit and inter-orbit coordination, respectively. The ablation analysis first assesses the intra-orbit coordination by comparing MPSLL with a “Symmetric” baseline that uses symmetric trees with all satellites participating in model distribution and aggregation. Fig. 5(e) shows the result: two solid curves at the top compare MPSLL and Symmetric, where the inter-orbit coordination is fixed to RFL. The result shows improvement by MPSLL. A similar result is observed from two dotted curves at the bottom, which compare MPSLL and Symmetric by fixing the inter-orbit coordination to FedFa. Next, we perform an ablation analysis of the inter-

orbit coordination by fixing the intra-orbit coordination to MPSLL and substituting RFL with other FL algorithms. The result in Fig. 5(f) reveals the consistent superiority of RFL. Specifically, RFL surpasses FedFa by alleviating the amnesia problem. Compared to FedAsyn (linear) and FedBuff, RFL is more effective in using local data, which promotes faster convergence.

V. DISCUSSION AND LIMITATION

The proposed system decouples satellite FL into intra-orbit and inter-orbit coordination to address distinct heterogeneities. Intra-orbit coordination dynamically selects participants based on system heterogeneity. Orbit-level replacement with staleness-aware aggregation ensures global knowledge persistence. These design choices offer deployment advantages. Since MPSLL executes at the GS, satellites only perform local training and model relay, respecting onboard compute limits. Intra-orbit ISLs reduce ground-satellite transmissions to two per iteration, which minimizes protocol overhead. As the GS serves as the central coordinator in standard satellite architectures, RFL requires no firmware modification and is compatible with existing infrastructure.

However, if there exist nodes with persistent heterogeneity, it may lead to unequal selection priorities and imbalanced participation across iterations. For example, for a satellite with

very poor hardware and ISL links, its participation in FL may be relatively limited. In future work, we will explore improving the fairness of our method by using historical participation records to prioritize underrepresented satellites.

VI. CONCLUSIONS

This article proposes a heterogeneity-aware satellite FL system that explicitly accounts for the heterogeneity in communication and computation latencies. We introduce the MPSLL algorithm for adaptive intra-orbit coordination and RFL for fully asynchronous inter-orbit coordination. The proposed joint design effectively mitigates the negative impact of system heterogeneity and significantly improves the communication efficiency of satellite FL in dynamic satellite environments.

ACKNOWLEDGMENT

This work was supported in part by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (IITP-2026-RS-2020-II201819, ICT Creative Consilience Program, 10%, and No. RS-2026-25507543, Development of AI Co-Scientist based on Scientific Causal World Model, 70%, and No. 2021-0-01715, Development of Communications Payload for GEOKOMPSAT-3); the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2022-NR070834, RS-2026-25482946); the Zhejiang Provincial Natural Science Foundation of China under Grants LZ25F010003 and LQN26F020066; the National Natural Science Foundation of China under Grant 62401506; the Zhejiang Provincial Key Research and Development Program under Grant 2026C02A1244; the Central Government Guiding Funds for Local Science and Technology Development under Grants 2025ZY01069 and 2026ZY01018; the Fundamental Research Funds for the Provincial Universities of Zhejiang under Grant FR2501ZD.

REFERENCES

- [1] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine learning and systems*, vol. 2, pp. 429–450, 2020.
- [2] C. Xie, S. Koyejo, and I. Gupta, "Asynchronous federated optimization," *arXiv preprint arXiv:1903.03934*, 2019.
- [3] J. Nguyen, K. Malik, H. Zhan, A. Yousefpour, M. Rabbat, M. Malek, and D. Huba, "Federated learning with buffered asynchronous aggregation," in *Proc. Artif. Intell. Statist. (AISTATS)*. PMLR, 2022, pp. 3581–3607.
- [4] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "On-board federated learning for satellite clusters with inter-satellite links," *IEEE Trans. Commun.*, vol. 72, no. 6, pp. 3408–3424, 2024.
- [5] W. Liu, Z. Lai, Q. Wu, H. Li, Q. Zhang, Z. Li, Y. Li, and J. Liu, "In-orbit processing or not? sunlight-aware task scheduling for energy-efficient space edge computing networks," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Vancouver, BC, Canada, May 2024, pp. 881–890.
- [6] J. Lyu, Y. L. Guan, X. Liu, and E. Gunawan, "Inter-orbital inter-satellite communication link performance analysis at thz band considering gravity-induced antenna pointing error," *IEEE Trans. Veh. Technol.*, vol. 74, no. 6, pp. 10000–10005, 2025.
- [7] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "On-board federated learning for dense leo constellations," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2022, pp. 4715–4720.
- [8] M. Elmahallawy and T. Luo, "Optimizing federated learning in leo satellite constellations via intra-plane model propagation and sink satellite scheduling," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2023, pp. 3444–3449.

- [9] N. Razmi, B. Matthiesen, A. Dekorsy, and P. Popovski, "Energy-aware federated learning in satellite constellations," in *Proc. IEEE Globecom Wkshps.* IEEE, 2024, pp. 1–6.
- [10] Z. Lin, Z. Chen, Z. Fang, X. Chen, X. Wang, and Y. Gao, "Fedsn: A federated learning framework over heterogeneous leo satellite networks," *IEEE Trans. Mobile Comput.*, vol. 24, no. 3, pp. 1293–1307, 2025.
- [11] J. Lai, H. Liu, G. Xu, W. Jiang, X. Wang, and D. Jiang, "Joint computation offloading and resource allocation for leo satellite networks using hierarchical multi-agent reinforcement learning," *IEEE Trans. Cogn. Commun. Netw.*, vol. 11, no. 4, pp. 2554–2567, 2025.
- [12] B. Matthiesen, N. Razmi, I. Leyva-Mayorga, A. Dekorsy, and P. Popovski, "Federated learning in satellite constellations," *IEEE Netw.*, vol. 38, no. 2, pp. 232–239, 2024.
- [13] H. Xu, Z. Zhang, S. Di, B. Liu, K. A. Alharthi, and J. Cao, "Fedfa: A fully asynchronous training paradigm for federated learning," in *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, Aug. 2024, pp. 5281–5288.
- [14] P. Helber, B. Bischke, A. Dengel, and D. Borth, "Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 7, pp. 2217–2226, 2019.
- [15] I. Leyva-Mayorga, B. Soret, and P. Popovski, "Inter-plane inter-satellite connectivity in dense leo constellations," *IEEE Trans. Wireless Commun.*, vol. 20, no. 6, pp. 3430–3443, 2021.

BIOGRAPHIES

Bohang Jiang received the B.E. degree in network engineering and the M.E. degree in information and communication engineering from Zhejiang Gongshang University, Hangzhou, China, in 2021 and 2024, respectively. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, College of Informatics, Korea University, Seoul, Korea. His research interests include distributed optimization, federated learning, and satellite networks.

Chao Chen (Member, IEEE) received the B.S. degree in communication engineering from Harbin Institute of Technology, Harbin, China, in 2009, and the Ph.D. degree in electrical engineering from Korea University, Seoul, Korea, in 2015. Since 2015, he has been with the School of Information and Electronic Engineering (Sussex Artificial Intelligence Institute), Zhejiang Gongshang University, Hangzhou, China, where he is currently an Associate Professor. His research interests include wireless communication, space-air-ground integrated network, and mobile edge computing.

In-Sop Cho received the B.S. and Ph.D. degrees in radio communications engineering from Korea University, Seoul, Korea, in 2011 and 2021, respectively. In 2021, he was with the Korea Electric Power Research Institute as a senior researcher. Since 2022, he has been with the Electronics and Telecommunications Research Institute (ETRI), Daejeon, Korea, where he is currently a senior researcher. His research interests include resource optimization for wireless communication, network scheduling, and satellite communication.

Min-Su Shin received the B.S. and M.S. degrees in electronics engineering from Korea Aerospace University, Goyang, Korea, in 1998 and 2000, respectively, and the Ph.D. degree in information communication engineering from Chungnam National University, Daejeon, Korea, in 2011. He joined the Electronics and Telecommunications Research Institute (ETRI), Daejeon, Korea, in 2000, where he has been involved in satellite communication system development. His current research interests include low Earth orbit (LEO) satellite network design, specifically LEO network routing mechanism and operation.

Inkyu Lee (Fellow, IEEE) received the B.S. degree (Hons.) in control and instrumentation engineering from Seoul National University, Seoul, Korea, in 1990, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1992 and 1995, respectively. Since 2002, he has been with Korea University, Seoul, Korea, where he is currently a Professor with the School of Electrical Engineering. His research interests include digital communications and signal processing techniques applied to next-generation wireless systems.

Seung Jun Baek (Member, IEEE) received the B.S. degree from Seoul National University, Seoul, Korea, in 1998, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Texas at Austin, Austin, TX, USA, in 2002 and 2007, respectively. From 2007 to 2009, he was a Member of Technical Staff with the DSP Systems R&D Center, Texas Instruments. In 2009, he joined the College of Informatics, Korea University, Seoul, Korea, where he is currently a Professor. His research interests include deep learning, data-driven optimization, mobile computing, and fog/edge computing.